

CLAIMS

What is claimed is:

- 1 1. A method for processing a database statement within a database server, the method
2 comprising the steps of:
3 receiving at the database server the database statement;
4 determining that at least one operation required by the database statement can be
5 parallelized;
6 within the database server, generating a set of information about how to execute the
7 database statement;
8 causing a plurality of slave processes to perform said at least one operation by
9 sharing the set of information with each slave process of said plurality of slave
10 processes, wherein the set of information shared with each slave
11 process includes
12 (a) information about a task to be performed by said slave process, and
13 (b) information about one or more tasks, to be performed by processes
14 other than the slave process, to execute the database statement; and
15 sending to each slave process of said plurality of slave processes data that
16 indicates which part of the at least one operation should be performed
17 by the slave process.
- 1 2. The method of Claim 1 wherein:
2 the step of sharing the set of information includes sharing an execution plan for the
3 database statement; and

4 sharing the execution plan with a particular slave process of the plurality of slave
5 processes is performed by:
6 providing an original statement of the database statement to a node on which
7 the particular slave process resides, wherein the original statement is
8 the form of the database statement in which the database statement
9 was received by the database server;
10 at said node, generating an equivalent execution plan based on the original
11 statement; and
12 the particular slave process accessing the equivalent execution plan.

1 3. The method of Claim 2 wherein:
2 further comprising the step of providing to the node additional information that
3 includes at least one of (a) values associated with session parameters of a
4 database session in which the database statement was received, and (b) values
5 associated with optimizer parameters that were used by an optimizer to
6 generate a plan for the database statement in a node other than said node; and
7 the step of generating an equivalent execution plan is performed based, at least in
8 part, on the additional information.

1 4. The method of Claim 1 wherein:
2 the step of generating a set of information includes generating an execution plan for
3 the database statement, wherein the set of information includes the execution
4 plan; and
5 the step of sending to each slave process of said plurality of slave processes data that
6 indicates which part of the at least one operation should be performed by the

7 slave process includes sending to each slave process data that indicates a
8 specific portion of the execution plan that is to be performed by the slave
9 process.

1 5. The method of Claim 4 wherein:

2 the step of sending to each slave process data that indicates a specific portion of the
3 execution plan that is to be performed by the slave process includes sending to
4 a particular slave process data that indicates a particular portion of the
5 execution plan that is to be performed by the particular process; and
6 the method further includes the step of the particular slave process determining how
7 to execute the particular portion based, at least in part, on characteristics of the
8 execution plan other than the particular portion of the plan that is to be
9 executed by the particular slave process.

1 6. The method of Claim 1 wherein:

2 the step of generating a set of information includes
3 generating an execution plan for the database statement;
4 constructing a shared cursor for the database statement, wherein the shared
5 cursor provides access to the execution plan; and
6 the step of sharing access includes providing each slave process of said plurality of
7 slave processes access to the shared cursor.

1 7. The method of Claim 6 wherein the step of providing each slave process of said
2 plurality of slave processes access to the shared cursor includes allowing two or more of said
3 slave processes to access a shared instance of the shared cursor.

1 8. The method of Claim 6 wherein the step of providing each slave process of said
2 plurality of slave processes access to the shared cursor includes allowing one of the slave
3 processes to access a first instance of the shared cursor, and allowing another one of the slave
4 processes to access a second instance of the shared cursor.

1 9. The method of Claim 8 wherein:
2 the one slave process resides on a first node;
3 the other slave process resides on a second node; and
4 the first node is a different node than said second node.

1 10. The method of Claim 9 wherein:
2 a first plurality of slave processes on said first node share access to said first instance
3 of said shared cursor; and
4 a second plurality of slave processes on said second node share access to said second
5 instance of said shared cursor.

1 12. The method of Claim 1 wherein:
2 the step of generating a set of information includes generating an execution plan for
3 the database statement, wherein the set of information includes the execution
4 plan; and
5 the method further comprises the step of inserting into the execution plan a granule
6 iterator row source that encapsulates a horizontal partitioning of a base object
7 upon which the database statement operates.

1 13. The method of Claim 1 wherein:

2 the step of generating a set of information includes generating an execution plan for
3 the database statement, wherein the set of information includes the execution
4 plan; and
5 the method further comprises the step of inserting into the execution plan at least one
6 distribution row source that specifies how data is to be redistributed between
7 one of
8 a first slave set and a query coordinator; and
9 a first slave set and a second slave set.

1 14. The method of Claim 13 wherein the step of inserting into the execution plan at least
2 one distribution row source includes:

3 inserting into the execution plan at least one sender-side distribution row source that
4 indicates how sending processes are to distribute data that the sending
5 processes produce; and
6 inserting into the execution plan at least one receiver-side distribution row source that
7 indicates how receiving processes are to obtain data that the receiving
8 processes are to consume.

1 15. The method of Claim 1 wherein:

2 the step of generating a set of information includes generating an execution plan for
3 the database statement, wherein the set of information includes the execution
4 plan; and
5 the method further comprises the step of inserting into the execution plan a
6 parallelizer row source that encapsulates the scheduling of tasks that slave
7 processes are to perform.

1 16. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 1.

1 17. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 2.

1 18. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 3.

1 19. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 4.

1 20. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 5.

1 21. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 6.

1 22. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 7.

1 23. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 8.

1 24. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 9.

1 25. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 10.

1 26. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 11.

1 27. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 12.

1 28. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 13.

1 29. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 14.

1 30. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to perform the
3 method recited in Claim 15.

4